

CLAIMS

What is claimed is:

1. A method for classifying an incoming packet, comprising:
 - maintaining a database associated with one or more patterns of fields, wherein the database is developed by:
 - mapping each pattern to a unique numeric identifier, wherein the number of unique numeric identifiers is equal to the number of patterns, and wherein a size of each unique numeric identifier is substantially smaller than the field of each pattern; and
 - determining a range of one or more of the unique numeric identifiers to be associated with each pattern, wherein the range is bounded by a minimum unique numeric identifier and a maximum unique numeric identifier;
 - using a field of the incoming packet to determine an associated identifier for that field, wherein the associated identifier is equal to one of the unique numeric identifiers;
 - matching the associated identifier assigned to the field of the incoming packet with one or more of the ranges; and
 - determining how to process the incoming packet based on the act of matching.
2. The method of claim 1, wherein the fields of the patterns and the fields of the incoming packets are network addresses.
3. The method of claim 1, wherein the incoming packet has more than one field, and wherein the act of using a field of the incoming packet to determine an associated identifier is carried out for each of the fields of the incoming packet.
4. The method of claim 1, further comprising, if the associated identifier matches one of the ranges bounded by different unique numeric identifiers, determining priority among the unique numeric identifiers in the range and selecting the highest priority for processing.
5. The method of claim 1, wherein the database is a trie.
6. The method of claim 5, wherein the size of each numeric identifier is equal to the largest integer greater than $\log_2 N$, where N is the number of unique patterns.
7. The method of claim 5, wherein the trie is developed using a depth-first, binary traversal of the patterns.
8. The method of claim 1, wherein using a field of the incoming packet to determine an associated identifier includes using a longest-prefix matching (LPM) process.

9. The method of claim 1, further comprising:
 - maintaining a second database associated with the one or more patterns of fields; and
 - exchanging the second database for the database.
10. The method of claim 1, wherein using a field of the incoming packet to determine an associated identifier includes looking up the associated identifier using addressable memory.
11. The method of claim 1, wherein the database is developed by constructing a trie, wherein the leaves of the trie are developed by:
 - selecting a highest bit to correspond to a root leaf; and
 - parsing through one of the fields of one of the patterns from left to right to develop child leaves.
12. A method for developing a database of patterns of fields for use in packet classification, comprising:
 - mapping each pattern to a unique numeric identifier, wherein the number of unique numeric identifiers is equal to the number of patterns, and wherein a size of each unique numeric identifier is substantially smaller than the field of each pattern; and
 - determining a range of one or more of the unique numeric identifiers to be associated with each pattern, wherein the range is bounded by a minimum unique numeric identifier and a maximum unique numeric identifier.
13. The method of claim 12, further comprising building a trie based on the fields of the patterns to develop the database.
14. The method of claim 13, wherein mapping each pattern includes associating one of the unique numeric identifiers to each terminal node of the trie.
15. An apparatus for classifying an incoming packet, comprising:
 - a database associated with one or more patterns of fields, wherein the database is developed by:
 - mapping each pattern to a unique numeric identifier, wherein the number of unique numeric identifiers is equal to the number of patterns, and wherein a size of each unique numeric identifier is substantially smaller than the field of each pattern; and
 - determining a range of one or more of the unique numeric identifiers to be associated with each pattern, wherein the range is bounded by a minimum unique numeric identifier and a maximum unique numeric identifier;

means for using a field of the incoming packet to determine an associated identifier for that field, wherein the associated identifier is equal to one of the unique numeric identifiers;

means for matching the associated identifier assigned to the field of the incoming packet with one or more of the ranges; and

means for determining how to process the incoming packet based on the act of matching.

16. A system for packet classification, comprising:

a pre-processor to receive a packet header of an incoming packet, the packet header including a field, and to assign an associated identifier to the field;

a first memory device, the first memory device including a first set of binary patterns;

a second memory device, the second memory device including a second set binary patterns;

a third memory device, the third memory device including instructions for applying one of the first and second sets of binary patterns to the associated identifier assigned to the field of the incoming packet; and

a processor to apply the instructions to the field to match the field to one of the patterns in the first set of binary patterns or second set of binary patterns.

17. The system of claim 16, wherein the field is a network address.

18. The system of claim 26, further comprising:

a fourth memory device, the fourth memory device including a result corresponding to a matched pattern received from the processor.

19. The system of claim 26, further comprising:

means for determining priority among one or more matched patterns.

20. A method for packet classification, comprising:

receiving a packet header of an incoming packet, the packet header including a field;

assigning an associated identifier to the field;

maintaining a first set of binary patterns and a second set binary patterns;

applying one of the first and second sets of binary patterns to the associated identifier assigned to the field of the incoming packet; and

matching the field of the incoming packet to one of the patterns in the first set of binary patterns or second set of binary patterns.

21. An apparatus for packet classification, comprising:
- (a) a processor array containing a plurality of processing elements, wherein each processing element is configured to: (i) compare a classification record derived from an incoming packet with one or more patterns associated with rules to be enforced by performing programmed relational operations, and (ii) capture state information in a shift register that is shared by more than one of the processing elements;
 - (b) an instruction decoder to suspend operations for a set of the processing elements and to restart operations for the set of processing elements;
 - (c) at least one priority encoder to determine a highest priority from the state information from the shift registers; and
 - (d) a record memory that is addressed for read and write capabilities, wherein writes are executed to indicate which of several simultaneously processed packets is being submitted to the apparatus for packet classification, and wherein reads are executed to indicate which bit or bits are to be provided to the processor array from all of the simultaneously processed packets.
22. The apparatus of claim 21, wherein the programmed relational operations include equality, less than or equal to, and greater than or equal to operations.
23. The apparatus of claim 21, further comprising an instruction memory containing instructions to: (i) select which bit or bits of the simultaneously processed packets are to be examined by the processor array, and (ii) determine how the selected bit or bits are to be compared against one or more of the patterns.
24. The apparatus of claim 23, wherein the instructions in the instruction memory are a classification program, wherein the classification program contains further instructions to divide the classification program into subsets of sequential instructions, each of which represents a subset of the rules defined by the classification program.
25. The apparatus of claim 24, wherein a prioritized relationship exists between a pair of the subsets of the rules.
26. The apparatus of claim 24, wherein the record memory is rewritable during execution of the classification program such that a result from the priority encoder can be subsequently examined by the classification program.
27. The apparatus of claim 26, wherein the result includes the current highest priority match as determined by the priority encoder.

28. The apparatus of claim 24, wherein one or more bits selected from one of the classification records in the record memory can determine a subset of instructions in the classification program to apply to one or more of the classification records.
29. The apparatus of claim 28, wherein the one or more bits selected are selected under programmed control.
30. The apparatus of claim 28, wherein the classification program can be seamlessly updated.
31. The apparatus of claim 28, wherein the effective size of the classification program can be increased.
32. The apparatus of claim 28, wherein the number of processing elements applied to a single incoming packet can be scaled.
33. The apparatus of claim 28, wherein the number of simultaneously processed packets can be scaled.
34. The apparatus of claim 28, wherein the number of packet bits used to select a classification program can be scaled.
35. The apparatus of claim 28, wherein multiple instances of the apparatus are cascaded such that a first instance of the apparatus feeds into a second instance of the apparatus.
36. The apparatus of claim 21, wherein multiple instances of the apparatus are cascaded such that a first instance of the apparatus feeds into a second instance of the apparatus.